XR for Complex Datacenter Environments

Why XR Needs AAS and AI-Driven Interaction to Support Operators in Next-Generation Datacenters

Stephan Waltermann[1,*,†], Andre Borrmann[1,†], Michel Munzert[1,†], Frank Feinbube[1,*,†] and Bosen Zhou[1,*,†]

[1] *SAP SE*

Abstract

Modern datacenters are officially classified as critical infrastructure and are transforming into highly complex, AI-driven environments characterized by rising energy demands, heterogeneous assets, and strict operational constraints. As these facilities evolve, they increasingly resemble challenging environments in which operators must manage dense, noisy, high-risk, and often partially inaccessible spaces. Extended Reality (XR) offers promising benefits for spatial awareness, guided procedures, and remote collaboration in such contexts, yet existing deployments in industrial settings reveal persistent limitations related to data consistency, workflow integration, and operator cognitive load. This position paper argues that XR alone cannot deliver reliable operational value in mission-critical datacenter settings without a robust semantic and adaptive foundation. We propose the combination of the Asset Administration Shell (AAS) as a unified and machine-interpretable data layer, XR as an immersive visualization medium, and AI-assisted interaction ("AI for UX") as an adaptive support mechanism. Together, these components form an integrated architecture that enhances data quality, reduces cognitive burden, improves situational awareness, and enables future scenarios such as XR-mediated teleoperation in sealed or hazardous datacenter environments. We outline the opportunities of this approach, highlight remaining challenges, and identify open research questions essential for deploying XR in demanding operational contexts.

Keywords

Virtual Reality (VR), Extended Reality (XR), Asset Administration Shell (AAS), Datacenter, Critical infrastructure, Challenging environments

1. Introduction

Modern datacenters have become some of the most complex and mission-critical infrastructures in today's digital ecosystem. Their rapid growth in energy consumption—already estimated at 1–2% of global electricity use [1] —and the increasing dominance of AI-centric workloads place unprecedented pressure on operators. At the same time, global data generation is projected to reach 463 exabytes per day by 2025 [2], further intensifying demands for precise monitoring, troubleshooting, and coordination across heterogeneous systems. These developments make datacenter environments highly technical, densely instrumented, and increasingly difficult to manage through traditional dashboards or isolated monitoring tools.

Extended Reality (XR), while well-established in domains such as industrial training and knowledge work [3, 4], has only recently been explored for datacenter operations . XR offers unique advantages in such environments: spatially aligned representations of racks and cable routes, visualization of otherwise invisible parameters like airflow or thermal conditions, and intuitive guidance during maintenance tasks. These capabilities are particularly valuable in high-density or high-risk environments that challenge situational awareness. [5]

However, XR adoption in mission-critical infrastructures is still limited. Studies on XR in demanding operational settings highlight persistent concerns related to reliability, environmental constraints,

CHI2026: XR4CE: CHI Workshop on XR for Challenging Environments, April 13–17, 2026, Barcelona, Spain

*Corresponding author.

†These authors contributed equally.

stephan.waltermann@sap.com (S. Waltermann); a.borrmann@sap.com (A. Borrmann); michel.munzert@sap.com (M. Munzert); frank.feinbube@sap.com (F. Feinbube); bosen.zhou@sap.com (B. Zhou)

© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

safety, and device robustness [6]. These concerns mirror challenges in datacenters, where inaccurate, incomplete, or inconsistent data can make immersive interfaces misleading rather than helpful. Likewise, existing XR systems struggle with interoperability across heterogeneous assets and lack adaptive interfaces capable of supporting operators under high cognitive load. [7]

This position paper argues that XR alone cannot deliver reliable operational value in datacenters without stronger data and interaction foundations. We propose combining the Asset Administration Shell (AAS) as a unified semantic layer with AI-assisted interaction ("AI for UX") and XR-based visualization. Together, these elements establish a resilient ecosystem in which (1) datacenter information is harmonized and machine-interpretable, (2) XR interfaces can draw on trustworthy and consistent data, and (3) AI provides adaptive guidance that reduces cognitive demand. Beyond technical integration, the approach is grounded in human-factors considerations: it aims to strengthen operators' situational awareness, sensemaking, and ability to manage time-critical tasks within increasingly complex infrastructure landscapes.

2. Context and Motivation

2.1. Changing Landscape of Datacenter Operations

AI-centric workloads are reshaping datacenter operations, driving rapid growth in computational demand, power consumption, and system complexity. Recent studies show that large-scale AI model training imposes unprecedented electrical loads on datacenters, significantly increasing pressure on energy infrastructures and operational planning. At the same time, datacenters face new requirements for ultra-low latency, higher energy efficiency, and more automated management as AI becomes a dominant workload category. The shift toward new rack architectures and next-generation cooling systems for GPU-dense workloads is creating immediate pressure to remodel existing datacenter layouts and physical infrastructure. High-density GPU clusters demand redesigned racks, higher power delivery, and additional spacing or cooling capacity, leading many facilities to undergo retrofitting or even full architectural updates to support liquid-based or hybrid cooling. Because these technologies are still emerging and evolve quickly, operators face added challenges in adapting the physical layer, managing rising energy usage, and upgrading legacy environments to become AI-ready. These developments intersect with rising sovereignty constraints, where regulations and data-locality requirements affect workload placement and cross-regional orchestration. As a result, scheduling compute-intensive workloads across heterogeneous, multi-region environments becomes increasingly challenging, especially when balancing performance, compliance, and resource availability. [8, 9]

2.2. Opportunities and Uncertainties of XR in Datacenters

XR offers promising opportunities for improving spatial awareness, guided procedures, and remote collaboration in complex technical environments. Empirical evaluations in industrial settings demonstrate that XR-based interfaces can enhance task performance, deepen operators' understanding of system states, and improve acceptance when interacting with digital twins and machinery. In datacenter operations, XR can directly support routine activities such as checking rack inventory, identifying cabling issues, or planning infrastructure changes. These tasks require operators to combine information from multiple systems under time pressure, which increases cognitive load. XR-based digital twins can ease this burden by aligning infrastructure data with spatial context and improving situational understanding. However, several uncertainties remain regarding its suitability for datacenter operations. Research highlights challenges such as device fragmentation, usability problems, and the difficulty of integrating XR tools into existing organizational workflows, especially where accuracy, stability, and data consistency are critical requirements. In addition, stakeholders often express skepticism due to privacy concerns, workflow disruptions, and limited evidence of long-term reliability in mission-critical industrial environments. These open questions underline the need for robust data foundations and adaptive interaction mechanisms and clearly defined usage scenarios if XR is to become a sustainable component of future datacenter toolchains. [4, 10]

3. Problem Statement: Missing Integration Layer

While XR offers potential benefits for spatial understanding and guidance, studies show that its effectiveness in industrial environments depends heavily on complete, accurate, and consistently structured data. XR-based digital twin evaluations demonstrate performance degradation when underlying data is inconsistent or incomplete, limiting real-world applicability. In addition, industrial deployments reveal workflow disruptions and organizational barriers, indicating that XR tools do not integrate smoothly into established operational practices. These issues manifest differently across modalities: VR environments rely entirely on synthesized data, meaning errors often become visible only when they are severe or highly disruptive, whereas AR systems can misalign or display incorrect overlays even with minor inaccuracies in the underlying digital representations.

A key limitation is the absence of a unified semantic layer to harmonize heterogeneous datacenter assets. Research on distributed and heterogeneous computing environments shows that lack of standardization increases system complexity and impairs coordination across diverse resources. Modern AI-driven datacenters further amplify this challenge by requiring accurate, real-time data to support rapidly fluctuating workloads. [11] From an operational perspective, these inconsistencies also increase cognitive effort for operators. Infrastructure state, configuration, and monitoring data are distributed across multiple systems, requiring operators to reconcile dashboards, inventory records, and physical equipment during troubleshooting and maintenance tasks—often under significant time pressure. Moreover, XR interfaces often lack adaptive mechanisms needed to keep cognitive load manageable. Reviews of XR in industrial training show that insufficient adaptivity can elevate mental workload, reducing usability in demanding environments. [4]

In summary, XR cannot provide operational value in datacenters without (1) semantically consistent and complete data, (2) adaptive interfaces that manage cognitive demand, (3) integration into established workflows, and (4) mechanisms to coordinate heterogeneous assets—together forming the missing integration layer this paper addresses.

4. Proposed Approach: AAS, XR, and AI for UX as Enablers

4.1. AAS as the Core Semantic Data Foundation

The Asset Administration Shell (AAS) provides a standardized digital representation of all assets within an Industry-4.0 landscape, including datacenter components such as server rooms, racks, servers, and software modules. [12] This standardization is achieved through a common metamodel that defines how assets are described in terms of properties, capabilities, and services—independent of vendor, technology stack, or lifecycle phase. [13]

A central component of this model is the concept of submodels. Each submodel provides a domain-specific, reusable representation of an asset perspective— for example, identification, technical data, or maintenance. Defined through standardized templates, submodels enforce consistent structure and naming conventions, enabling plug-and-play interoperability: systems can consume and exchange information as long as they follow the AAS standard. To ensure a shared semantic understanding, the AAS references IEC 61360, defining how properties are described through standardized concept descriptions, value lists, and data types. [14] This semantic layer ensures that terms such as “temperature” are unambiguously understood across systems and domains. By combining standardized submodels with IEC 61360-based semantics, the AAS enables cross-vendor interoperability. This makes the AAS a neutral and vendor-independent carrier of asset information throughout the entire value chain. [15]

4.2. XR as the Adaptive Visualization Layer

While the AAS provides a structured and semantically consistent foundation for digital representations, the quantity and complexity of this information can quickly exceed what operators can efficiently interpret using conventional interfaces. XR technologies offer a modality that allows users to perceive complex relationships—such as spatial dependencies, structural hierarchies, or process flows—in a visual format closely aligned with real-world perception. [16] To remain effective in operational environments, immersive interfaces must also manage operator cognitive load and attentional demands. Techniques such as progressive disclosure or focus-plus-context visualizations enable users to shift from high-level datacenter overviews to detailed component diagnostics without losing spatial orientation or becoming overwhelmed by information.

XR also enables the visualization of otherwise invisible operational parameters, such as airflow, energy consumption, or thermal characteristics. By adapting the level of detail based on user context—such as proximity to components, gaze direction, or task type—XR can present only the most relevant information for a given activity. In this way, XR becomes a dynamic visualization layer that leverages AAS data as its semantic backbone.

A critical element of XR-based digital twins is the link between high-fidelity 3D assets and the structured, machine-interpretable representations provided by the AAS. Emerging standards such as OpenUSD offer promising capabilities for representing spatial information, anchor points, and interaction schemes. However, current ecosystems are fragmented: while platforms such as NVIDIA Omniverse support OpenUSD extensibility, many open-source rendering engines (e.g., Godot) and supporting libraries (e.g., Open3D) lack native support. Likewise, XR hardware ecosystems show inconsistent support for OpenXR frameworks, complicating cross-platform deployment and limiting broad adoption.

4.3. AI for UX as the Adaptive Interaction Layer

To complement XR-based visualization, AI can serve as an adaptive interaction layer that reduces cognitive load and supports operators in navigating complex datacenter environments. An AI companion embedded within the XR experience enables fully hands-free interaction by allowing users to communicate naturally—similar to conversing with a colleague—while the AI retrieves, filters, or explains relevant information. Because the AI has access to the underlying AAS data structures and the virtual environment, it can interpret context, anticipate user intentions, and deliver just-in-time assistance. This includes highlighting relevant components, summarizing system states, or guiding operators through multi-step procedures. In future scenarios, such an AI companion could also execute certain actions autonomously or semi-autonomously—such as triggering diagnostics, adjusting system parameters, or coordinating robotic agents—further extending operator capabilities in complex environments. To remain appropriate for operational environments, such AI systems must act as decision-support co-pilots rather than autonomous controllers. They should expose the underlying data sources and confidence levels of their recommendations, enabling operators to inspect reasoning and override actions at any time. This transparency supports calibrated trust and ensures that final decisions remain under human authority, particularly in cases where AI-generated insights conflict with observed system states. By combining XR's immersive spatial interface with AI-driven adaptivity, the interaction becomes more intuitive, efficient, and aligned with operator needs in high-complexity situations.

4.4. Synergy with Emerging Datacenter Trends

The combined use of AAS, XR, and AI creates a holistic, end-to-end operational view of modern datacenters. AAS ensures semantic consistency and interoperability across heterogeneous assets; XR provides intuitive, spatialized visualization of complex environments; and AI-supported interaction reduces cognitive burden and enhances task execution. This synergy also aligns with emerging trends in next-generation datacenters, such as increased sustainability requirements, operational automation, or environments that may no longer be safely accessible to humans. For example, future high-density or gas-cooled datacenter designs may require sealed environments, where XR-mediated teleoperation of robots becomes a practical necessity rather than a speculation. By building on AAS-based semantics and AI-driven guidance, such scenarios could enable operators to interact with critical infrastructure remotely, safely, and efficiently illustrating the long-term potential of the proposed architecture. [17]

5. Implications and Open Questions

The integration of AAS, XR, and AI has the potential to improve transparency and operator support in increasingly complex datacenter environments. However, several open questions remain. First, although XR can enhance situational awareness and guided procedures, its concrete value in everyday datacenter operations is still unclear, and operator acceptance in mission-critical contexts varies. Second, the approach depends heavily on the quality and completeness of AAS-based data ecosystems, which require consistent modeling practices and reliable lifecycle information across heterogeneous assets. Third, AI-driven interaction models must mature to provide dependable, context-aware assistance without introducing additional cognitive or operational risks. Finally, sovereignty, security, and compliance considerations may restrict how XR and AI can access or present datacenter information. These questions highlight the need for further research before large-scale deployment is feasible.

6. Conclusion

Datacenters are becoming more complex and increasingly shaped by AI-driven workloads, requiring new tools for transparency and operational support. XR alone cannot meet these demands without structured, semantically consistent data and adaptive interaction mechanisms. By combining the AAS as a shared semantic foundation with XR for spatial visualization and AI for context-aware interaction, we outline a cohesive approach to support next-generation datacenter operations. This architecture offers a promising path toward more intuitive workflows, improved situational awareness, and future scenarios in which operators interact with sealed or automated environments through XR-mediated robotic telepresence. Future work will examine how such systems can further strengthen operator sensemaking and decision-making in complex infrastructure landscapes.

Acknowledgments

- Acknowledgement: "This research work was carried out as part of the project "Apeiro Reference Architecture", in short ApeiroRA, under the funding ID 13IPC007, financed by the European Union - NextGenerationEU and funded by the German Federal Ministry for Economic Affairs and Energy."
- Disclaimer: "Financed by the European Union - NextGenerationEU. The views and opinions expressed are solely those of the author(s) and do not necessarily reflect the views of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them."

Declaration on Generative AI

During the preparation of this work, the author(s) used Microsoft Copilot solely for language-related improvements such as grammar correction, wording refinement, and stylistic suggestions.After using this tool, the author(s) carefully reviewed, validated, and edited all content and take full responsibility for the final publication.

References

[1] S. Virmani, Investigating the environmental sustainability of data centers, American Journal of Student Research 3 (2025) 10–18.

[2] T. Porcino, E. Clua, Challenges in data visualization with extended reality devices, in: Proceedings of the 23rd IFIP TC 14 International Conference, Manaus, 2024. IFIP TC 14.

[3] A. Queiroz, J. N. Bailenson, K. P. Blair, D. L. Schwartz, C. Thille, A. D. Wagner, Extended realities and the future of knowledge work: Opportunities, in: 2024 IEEE Conference on Virtual Reality and 3D User Interfaces Abstracts and Workshops (VRW), IEEE, Orlando, 2024.

[4] G. W. Scurati, A roadmap for xr integration in industry: challenges and opportunities, in: 2025 IEEE International Symposium on Mixed and Augmented Reality Adjunct (ISMAR-Adjunct), IEEE, 2025.

[5] R. Weinschenk, Extended reality offers new training capabilities, https://intelligence.uptimeinstitute.com/resource/extended-reality-offers-new-training-capabilities, 2024.

[6] J. van der Hooft, H. Amirpour, M. T. Vega, Y. Sanchez, R. Schatz, T. Schierl, C. Timmerer, A tutorial on immersive video delivery: From omnidirectional video to holography, IEEE Communications Surveys & Tutorials (2023) 1336–1375.

[7] I. F. Akyildiz, H. Guo, Wireless communication research challenges for extended reality (xr), ITU Journal on Future and Evolving Technologies (2022).

[8] X. Chen, X. Wang, A. Colacelli, M. Lee, L. Xie, Electricity Demand and Grid Impacts of AI Data Centers: Challenges and Prospects, Technical Report, arXiv, 2025.

[9] J. Yu, Implementation strategies and development trends of artificial intelligence in data center operations, in: Computing, Internet of Things and Data Analytics, 2025.

[10] A. Sergeeva, C. S. N. Ribalta, G. Lenzini, Identifying ethical challenges in xr implementations in the industrial domain: A case of off-highway machinery, in: CHIXR/2025/08, 2025.

[11] G. Grego, F. Nenna, L. Gamberini, Xr-based digital twin for industry 5.0: A usability and user experience evaluation, in: Extended Reality. XR Salento 2024, Springer, Cham, 2024.

[12] J. van der Hooft, H. Amirpour, M. T. Vega, Y. Sanchez, R. Schatz, T. Schierl, C. Timmerer, A tutorial on immersive video delivery: From omnidirectional video to holography, IEEE Access (2024) 84863–84877.

[13] Industrial Digital Twin Association (IDTA), Specification of the asset administration shell – part 1: Metamodel, https://industrialdigitaltwin.org/wp-content/uploads/2025/11/IDTA-01001-3-1-2_AAS-Specification_Part1_Metamodel.pdf, 2025.

[14] Industrial Digital Twin Association (IDTA), Aas specification – part 3a: Data specification – iec 61360, https://industrialdigitaltwin.org/en/wp-content/uploads/sites/2/2025/08/IDTA-01003-a-3-1-1_AAS-Specification_Part3a_DataSpecification.pdf, 2025.

[15] S. Cavalieri, M. G. Salafia, A model for predictive maintenance based on asset administration shell, in: Proceedings (venue unspecified), 2020. Venue details not specified.

[16] J. Zhao, B. E. Riecke, J. W. Kelly, J. Stefanucci, A. Klippel, Editorial: Human spatial perception, cognition, and behaviour in extended reality, Frontiers in Virtual Reality (2023).

[17] W. Leclerc, Immersion cooling safety in a data center, https://www.datacenterdynamics.com/en/opinions/immersion-cooling-safety-in-a-data-center/, 2024.